

Futures Prices in a Production Economy with Investment Constraints*

Leonid Kogan[†], Dmitry Livdan[‡], and Amir Yaron[§]

This Draft: March 29, 2005

Abstract

We document a new stylized fact regarding the term-structure of futures volatility. We show that the relationship between the volatility of futures prices and the slope of the term structure of prices is non-monotone and has a “V-shape”. This aspect of the data cannot be generated by basic models that emphasize storage while this fact is consistent with models that emphasize investment constraints or, more generally, time-varying supply-elasticity. We develop an equilibrium model in which futures prices are determined endogenously in a production economy in which investment is both irreversible and is capacity constrained. Investment constraints affect firms investment decisions, which in turn determine the dynamic properties of their output and consequently imply that the supply-elasticity of the commodity changes over time. Since demand shocks must be absorbed either by changes in prices, or by changes in supply, time-varying supply-elasticity results in time-varying volatility of futures prices. Calibrating this model, we show it is quantitatively consistent with the aforementioned “V-shape” relationship between the volatility of futures prices and the slope of the term-structure.

*We are grateful to Kerry Back, Pierre Collin-Dufresne, Francis Longstaff and Craig Pirrong, as well as seminar participants at Northwestern University, Texas A&M University, 2004 Western Finance Association meeting, 2004 Society of Economic Dynamics meeting, and 2004 European Econometric Society meeting for useful suggestions. We also thank Krishna Ramaswamy for providing us with the futures data. Financial support from the Rodney L. White Center for Financial Research at the Wharton school is gratefully acknowledged.

[†]Sloan School of Management, Massachusetts Institute of Technology and NBER. Cambridge, MA 02142. Phone: (617) 253-2289. Fax: (617) 258-6855. E-mail: lkogan@mit.edu.

[‡]Bauer School of Business, University of Houston. Houston, TX 77204, Phone:(713) 743-4813. Fax: (713) 743-4755. Email: dlivdan@uh.edu.

[§]The Wharton School, University of Pennsylvania and NBER. Philadelphia, PA 19104. Phone: (215)-898-1241. Fax: (215) 898-6200. E-mail: yaron@wharton.upenn.edu.

1 Introduction

In recent years commodity markets have experienced dramatic growth in trading volume, the variety of contracts, and the range of underlying commodities. There also has been a great demand for derivative instruments utilizing operational contingencies embedded in delivery contracts. For all these reasons there is a widespread interest in models for pricing and hedging commodity-linked contingent claims. Besides practical interest, commodities offer a rich variety of empirical properties, which make them strikingly different from stocks, bonds and other conventional financial assets. Notable properties of futures include, among others: (i) Commodity futures prices are often "backwardated" in that they decline with time-to-delivery, (ii) Spot and futures prices are mean reverting for many commodities, (iii) Commodity prices are strongly heteroscedastic and price volatility is correlated with the degree of backwardation, and (iv) Unlike financial assets, many commodities have pronounced seasonalities in both price levels and volatilities.

The 'theory of storage' of Kaldor (1939), Working (1948, 1949) and Telser (1958) has been the foundation of the theoretical explorations of futures/forward prices and convenience yields (value of the immediate ownership of the physical commodity). Based on this theory researchers have adopted two approaches to modelling commodity prices. The first approach is mainly statistical in nature and requires an exogenous specification of the 'convenience yield' process for a commodity (e.g., Brennan and Schwartz (1985), Brennan (1991), and Schwartz (1997)). The second strand of the literature derives the price processes endogenously in an equilibrium valuation framework with competitive storage (e.g., Williams and Wright (1991), Deaton and Laroque (1992, 1996), Routledge, Seppi, and Spatt (2000)). The appealing aspect of this approach is its ability to link the futures prices to the level of inventories and hence derive additional testable restrictions on the price processes.

From a theoretical perspective the models based on competitive storage ignore the production side of the economy, and consequently they suffer from an important limitation. Inventory dynamics have little if any impact on the long-run properties of commodity prices, which in such models are driven mostly by the exogenously specified demand process. In particular, prices in such models tend to mean revert too fast relative to what is observed

in the data (see Routledge et. al. (2000)), and more importantly the model can not address the rich term-structure dynamics of return volatility.

In this paper we document an important new stylized fact regarding the property of the term structure of volatility of futures prices. We demonstrate that the relationship between the volatility of futures prices and the slope of the forward curve (the basis) is non-monotone and has a “V-shape”. Specifically, conditionally on a negatively sloped term structure, the relation between the volatility of futures prices and the slope of the forward curve is negative. On the other hand, conditionally on a positively sloped term structure, the relation between the volatility and the basis is positive. This aspect of the data cannot be generated by basic models that emphasize storage, since such models imply a monotone relation between futures price volatility and the slope of the forward curve (see Routledge et. al. (2000)).

In light of the aforementioned stylized fact, we explore an alternative model characterizing the mechanism of futures price formation. Future prices are determined endogenously in an equilibrium production economy featuring constraints on investment, such as irreversibility. Investment constraints affect firms investment decisions, which in turn determine the dynamic properties of their output. In particular, because of the binding constraints on investment, supply-elasticity of the commodity changes over time. Since demand shocks must be absorbed either by changes in prices, or by changes in supply, time-varying supply-elasticity results in time-varying volatility of futures prices. In our calibration below we show that the model can also generate these patterns in a manner that is quantitatively similar to the data.

There exists very little theoretical work investigating the pricing of futures on commodities using production economy framework. Casassus, Collin-Dufresne and Routledge (2004) also analyze a model of equilibrium spot and futures oil prices in a general equilibrium production economy but with two goods. In their model production of the consumption good requires two inputs: the consumption good and oil. Oil is produced by wells whose flow rate is costly to adjust. Unlike our paper, however, their focus is on providing an appropriate specification for the joint dynamics of spot price, convenience yield and interest rates. Moreover, their model features only one constraint (irreversibility of investment) required for generating the stylized fact that we document here.

The rest of the paper is organized as follows. In Section 2 we describe our data analysis regarding future prices. Section 3 develops the model’s economy as well as presents the spot price and prices of financial assets derived in competitive equilibrium setting. In Section 4, we study quantitative implications of the model. Section 5 provides conclusions.

2 Empirical Analysis

We concentrate our empirical study on three commodities: crude oil, heating oil, and unleaded gasoline. Our data consists of daily futures prices for three contracts: NYMEX heating oil contract (HO) for the period from 1979 to 2000, NYMEX light sweet crude oil contract (CL) for the period from 1982 to 2000, and NYMEX unleaded New York harbor gasoline futures (HU) for the period from 1985 to 2000. Following previous work by Routledge et. al (2000), the data is sorted by contract horizon with the ‘one-month’ contract being the contract with the earliest delivery date, the ‘two-month’ contract having next earliest delivery date, etc.¹. We consider contracts with up to 12 months to delivery when contract ‘matures’ since liquidity and data availability is good for these horizons for all three contracts used in this study.² Since we are using daily data, our dataset is sufficiently large: it ranges from 2500 to 3500 data points across different contracts and maturities.

Instead of directly using futures prices, $P(t, T)$, we use daily percent changes, $R(t, T) = \frac{P(t, T) - P(t-1, T)}{P(t-1, T)}$. Percent price changes are not susceptible as much as price levels to seasonalities and trends, and therefore their volatility is more suitable for empirical analysis. We then proceed by constructing the term structure of the unconditional and conditional volatilities of daily percent changes on futures for all three contracts. In calculating conditional moments we classify the futures curve at each date t as to whether one trading day earlier the curve was in backwardation or contango (based on the shortest and third shortest maturity prices

¹In our data set, for each of the commodities, on any given calendar day there are several contracts available with different time to delivery measured in days. The difference in delivery times between these contracts is at least 32 days or more. We utilize the following procedure for converting delivery times to the monthly scale. For each contract we divide the number of days it has left to maturity by 30 (the average number of days in a month), and then round off the resultant. For days when contract with time to delivery of less than 15 days is traded, we add one “month” to the contract horizon obtained using the above procedure for all contracts traded on such days. The data is then sorted into bins based on the contract horizon measured in months.

²We refer to this time to delivery as time to maturity throughout the paper.

at that time). Figures 1 - 3 show these conditional and unconditional daily volatilities for futures percent changes for CL, HO, and HU contracts respectively. The behavior of crude oil (CL) contracts was studied by Routledge et. al. (2000). We find, as they did, that the volatility of futures prices is higher when the forward curve is in backwardation. This has been interpreted as evidence in favor of the standard storage theories, emphasizing the effect of inventory stock-outs on price volatility. Note, however, that the pattern is reversed for heating oil (HO) futures. There is no discernable relation for gasoline (HU) contracts. These basic observations suggests that evidence in favor of the standard storage theory is, at best, mixed.

Next, we study the patterns in volatility of futures price in more detail. Specifically, we estimate a functional relation between the futures price volatility and the one-day lagged slope of the forward curve. Following the definition of conditional sample moments, the time series of the slope of the forward curve is constructed as a logarithm of the ratio of the futures price of the third shortest maturity in months available on any day t , $P(t, 3)$, to the future price of the shortest maturity, $P(t, 1)$, available on the same day. We start by using lagged slope as the only explanatory variable for realized volatility

$$|R(t, T)| = \alpha_T + \beta_T \ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] + \varepsilon_T(t). \quad (1)$$

Note that the relation (1) can potentially yield different information than that contained in Figures 1–3, which were obtained by simply splitting the sample based on the slope of the forward curve, since we are now estimating a different functional form. The term structure of β_T as well as the corresponding t -statistics are shown for different commodities in Figures 4 through 6. We also report these results in Tables 1 through 3 for T equal to 1, 5, and 10 months. The negative sign of β_T for all times to maturity (most of them in the case of the HU contract) is a common feature across contracts. This result seems to be at odds with the relations shown in Figures 1 through 3, where volatility conditional on backwardation is mostly higher than the unconditional volatility. The apparent inconsistency becomes less puzzling in light of the intuition of the model we present below. In particular, our theoretical results motivate one to look for a non-monotone relationship between the volatility of future prices and the slope of their term structure. For that we decompose the lagged slope into

positive and negative parts and use them as separate explanatory variables (i.e., use a piecewise linear regression on the slope of the term structure),

$$|R(t, T)| = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] \right)^- + \varepsilon_T(t), \quad (2)$$

where $(X)^\pm$ denotes the positive (negative) part of X respectively. Figures 4 through 6 as well as Tables 1 through 3 illustrate our results. Both $\beta_{1,T}$ and $\beta_{2,T}$ are statistically and economically significant for most maturities and across contracts. More importantly, $\beta_{1,T}$ and $\beta_{2,T}$ differ in sign: $\beta_{1,T}$ are positive and $\beta_{2,T}$ are negative. Therefore, the relation between the volatility of futures prices and the slope of the term structure of prices is non-monotone and has a “V-shape”: conditional volatility declines as a function of the slope when the latter is negative, and increases when the latter is positive.

We perform several robustness checks. Using the square of daily price changes instead of the absolute value, i.e., estimating conditional variance instead of conditional volatility, leads to very similar conclusions. In most cases, both $\beta_{1,T}$ and $\beta_{2,T}$ remain statistically significant for CL and HO contracts, but only marginally significant for HU contract.

As a final robustness check we split our sample into pre- and post-Gulf war sub-samples. We perform the same analysis as in the case of the full sample on the post-Gulf war sub-sample. We find the same “V-shape” in the relationship between the volatility of futures prices and the slope of the term structure of prices for CL and HO contracts. The “V-shape” disappears in the post-Gulf war sub-sample in the case of HU contracts.

3 The Model

In this section we present our model for spot prices and derive futures prices.

3.1 Setup

We consider a continuous-time infinite-horizon economy. We focus on a competitive industry populated by firms using identical production technology. Production in the economy is done by means of a production function that exhibits constant returns to scale

$$Q_t = XK_t, \quad (3)$$

where K_t is capital and X is the productivity of capital assumed to be constant. For the sake of simplicity we will assume below that $X = 1$. Our results can be easily adjusted to accommodate the case when X is a stochastic process. Firms can adjust their capital stock according to

$$dK_t = (i_t - \delta)K_t dt, \quad (4)$$

where i_t is the investment rate and δ is the capital depreciation assumed to be a nonnegative constant. We assume that investment is irreversible, i.e., $i_t \geq 0$, and the rate of investment is bounded, i.e., $i_t \in [0, \bar{i}]$. These investment frictions give rise to the dynamic properties of futures prices investigated below.

Firms sell their output in the spot market at price S_t . We assume that financial markets are complete and the firms' objective is to maximize their market value, which in turn is given

$$V_0 = E_0 \left[\int_0^\infty e^{-r_f t} S_t Q_t dt \right] \quad (5)$$

We assume that the expected value is computed under the risk-neutral measure and the risk-free rate r_f is constant.

The consumers in the economy are represented by the demand curve

$$Q_t = Y_t^{-1} S_t^{-\frac{1}{\gamma}}, \quad Q_t \in (0, \infty) \quad (6)$$

where unexpected changes in Y_t represent demand shocks. We assume that Y_t follows a geometric Brownian motion process

$$\frac{dY_t}{Y_t} = \mu_Y dt + \sigma_Y dW_t. \quad (7)$$

We also assume that $\gamma > 1$. Results for the case of $\gamma \leq 1$ are analogous and are available from the authors.

Since our model is formulated under the risk-neutral probability measure, while the empirical observations are made under the “physical” probability measure, one has to make an explicit assumption about the relation between these two measures, i.e., about the risk premium associated with the shock process dW_t . To keep our specification as simple as possible, we assume that the risk premium is constant, i.e., the drift of the demand shock Y_t under the “physical” probability measure is equal to $\mu_Y + \lambda$, where λ is an additional parameter of the model.

3.2 Equilibrium Investment and Prices

Following Lucas and Prescott (1971), we characterize the equilibrium investment policy as a solution of the social planner's problem. Specifically, we define the surplus function³

$$U(Y_t, Q_t) = \int_1^{Q_t} S_t(q) dq = Y_t^{-\gamma} \int_1^{Q_t} Q^{-\gamma} dQ = Y_t^{-\gamma} \frac{Q_t^{1-\gamma} - 1}{1-\gamma}. \quad (8)$$

The social planner maximizes total surplus net of investment costs

$$\max_{i_t \in [0, \bar{i}]} E_0 \left[\int_0^\infty e^{-r_f t} \left(Y_t^{-\gamma} \frac{K_t^{1-\gamma}}{1-\gamma} - i_t K_t \right) dt \right] \quad (9)$$

subject to the dynamics of the demand curve and the capital accumulation rule

$$\begin{aligned} dK_t &= (i_t - \delta) K_t dt, \\ \frac{dY_t}{Y_t} &= \mu_Y dt + \sigma_Y dW_t, \\ K_t &\geq 0, \quad \forall t \geq 0. \end{aligned}$$

The details of the solution are given in the Appendix. The equilibrium investment policy is given by

$$i_t^* = \begin{cases} \bar{i}, & \omega_t \leq \omega^* \\ 0, & \omega_t > \omega^* \end{cases}$$

where $\omega_t = \ln(K_t Y_t)$ and ω^* is given in the Appendix. The evolution of the state variable ω_t under the risk-neutral measure is governed by

$$d\omega_t = \left[-\delta + i_t^* + \mu_Y - \frac{1}{2} \sigma_Y^2 \right] dt + \sigma_Y dW_t. \quad (10)$$

Therefore, as long as $0 < \delta - \mu_Y + \frac{1}{2} \sigma_Y^2 < \bar{i}$, ω_t has a stationary long-run distribution

$$p(\omega) = \frac{2\mu^+ \mu^-}{\bar{i} \sigma_Y^2} \begin{cases} e^{-\frac{2\mu^+}{\sigma_Y^2}(\omega^* - \omega)} & \omega \leq \omega^* \\ e^{-\frac{2\mu^-}{\sigma_Y^2}(\omega - \omega^*)} & \omega > \omega^* \end{cases}. \quad (11)$$

where we define the two constants $\mu^- = \delta - \mu_Y + \frac{1}{2} \sigma_Y^2$ and $\mu^+ = \bar{i} - \mu^-$. The condition for stationarity of ω_t can be re-stated as $0 < \mu^- < \bar{i}$. The details of the derivation are provided

³Starting the integration at 1 is inconsequential for our analysis. Recall that Q is restricted to be strictly positive guaranteeing a well defined objective function in equation (9). Starting the integration at any positive point below one would only add a constant to the objective function in equation (9) not affecting the first order conditions.

in the Appendix. The spot price process of firms' output is related to the equilibrium capital stock by the inverse demand curve

$$S_t = [K_t Y_t]^{-\gamma} = e^{-\gamma \omega_t}, \quad (12)$$

It is worth mapping our general investment constraint model to the oil industry. Oil (Q) is the output produced using physical capital K (e.g., refineries, pipes). Implicitly we are assuming there is an infinite supply of underground oil, and production is constrained by the existing capital stock K . This supply of capital and consequently of oil-output leads to price fluctuations in response to demand shocks. Futures prices (volatility) respond depending on anticipated future production which depend on the degree to which investment is constrained.

3.3 Futures Prices

The futures contract is a claim on the good which is sold on the spot market at prevailing spot price S_t . The futures price is computed as the conditional expectation of the spot price under the risk-neutral measure:

$$P(t, T) = E_t[S_{t+T}] = E_t[e^{-\gamma \omega_{t+T}}], \quad \forall T \geq 0. \quad (13)$$

where $P(t, T)$ denotes the price of a futures contract at time t with maturity date $t + T$. Since Equation (13) implies that $P(t, T)$ is a local martingale, the conditional expectation of its instantaneous change should be equal to zero

$$E_t[dP(t, T)] = \left[-\frac{\partial P}{\partial T} + \left(-\delta + i^*(\omega) + \mu_Y - \frac{1}{2}\sigma_Y^2 \right) \frac{\partial P}{\partial \omega} + \frac{1}{2}\sigma_Y^2 \frac{\partial^2 P}{\partial \omega^2} \right] dT = 0. \quad (14)$$

As a result, $P(t, T)$ can be equivalently characterized as a solution to the above partial differential equation with a terminal condition

$$P(t, 0) = e^{-\gamma \omega_t}.$$

Since no analytical solution exists for this equation, we solve it numerically using a finite-difference scheme.

4 Estimation and Numerical Simulation

In this section we study how well our model can replicate quantitatively the key features of the behavior of futures prices reported in Section 1. We first estimate the model's parameters using a simulated method of moments. Our procedure is similar in spirit, but somewhat different technically, from those proposed in Lee and Ingram (1991) and Duffie and Singleton (1993). We then discuss additional implications of the model.

4.1 Simulated Moments Parameter Estimation

Estimation Procedure

Our goal is to estimate a vector of structural parameters, $\theta \equiv \{\gamma, \mu_Y, \sigma_Y, \bar{i}, r_f, \delta, \lambda\}$. We do this using a classical minimum distance (CMD) method, which requires matching a set of functions of structural parameters with the corresponding set of empirical estimates. Our procedure can be equivalently viewed as a generalized method of moments (GMM), since all the functions we consider can be expressed as sample moments. Let x_t be the vector-valued process of historical futures prices and output and consider a function of the observed sample $F_T(x)$, where T is the sample length. The statistic $F_T(x)$ could represent a collection of sample moments or even a more complicated estimator, such as the slope coefficients in a regression of volatility on the term structure as in (1). Assume that as the sample size T increases, $F_T(x)$ converges in probability to a limit $M(\theta)$, which is a function of structural parameters. Since many of the useful population moments cannot be computed analytically, we estimate them using Monte Carlo simulation. In particular, let $m_S(\theta) = \frac{1}{S} \sum_{s=1}^S F_T(x^s; \theta)$ represent the estimate of $M(\theta)$ based on S independent model based statistics, where x^s represents a vector valued process of simulated futures prices and output of length T based on simulating the model at parameter values, θ .⁴ Let $G_S(x, \theta) = m_S(\theta) - F_T(x)$, denote the difference between the estimated theoretical mean of the statistic F and its observed (empirical) value. Under appropriate regularity conditions, it can be shown that as the size of the sample, T , and the number of simulations S increase to infinity, the CMD (GMM)

⁴Specifically, for any given value of θ , we draw S realizations of the state variable ω_t from its long-run steady-state distribution (which itself depends on model parameters and is given by (11)). Then, for each set of initial conditions, we simulate a path of the state variable of the same length as the historical sample and evaluate the function $F(x, \theta)$ for each simulated path of the economy.

estimate of θ ,

$$\theta_S = \arg \min_{\theta} J_T = \arg \min_{\theta} G_S(x, \theta)' W_T G_S(x, \theta)$$

will be a consistent estimator of θ . The matrix W_T in the above expression is positive definite and assumed to converge in probability to a deterministic positive definite matrix W .

Assume that V is the asymptotic variance-covariance matrix of $F_T(x; \theta)$, i.e., $S^{-1/2} \sum_{s=1}^S F_T(x^s; \theta) \rightarrow N(0, V)$. Then, if we use the efficient choice of the weighting matrix, $W = V^{-1}$, the estimator θ_S is asymptotically normal,

$$S^{1/2}(\theta_S - \theta) \xrightarrow{d} N(0, (D'V^{-1}D)^{-1}), \quad (15)$$

where $D = \nabla_{\theta} M(\theta)$.

We perform estimation in two stages. During the first stage, we use an identity matrix for the weighting matrix W . During the second stage, the weighting matrix is set equal to the inverse of the estimated covariance matrix: $W = V_S^{-1}$, where $V_S = S^{-1/2} \sum_{s=1}^S F_T(x^s; \theta)$. To compute standard errors, we replace D in (15) with its consistent estimate $D_S = \nabla_{\theta} m_S(\theta)$.

We estimate the vector of seven model parameters, θ , by matching the unconditional properties of futures prices, specifically, the unconditional mean and volatility of daily percent price changes for futures of various maturities as well as the mean, volatility, and the 30-day autoregressive coefficient of the slope of the forward curve. We use crude oil futures prices as a benchmark. Reproducing unconditional properties of futures prices with a simple single-factor model is a nontrivial task, as we discuss below. However, the most novel implications of our model have to do with the conditional nature of the relation between the volatility of futures prices and the slope of the term structure. With this in mind, we choose model parameters to match the unconditional level of futures price volatility and then evaluate the quality of model predictions based on the conditional moments, which were not used in estimation.

Identification

Not all of the model parameters can be independently identified from the data we are considering. In this subsection we discuss the relations between structural parameters and observable properties of our model economy, which suggest which of the structural

parameters can be identified and what dimensions of empirical data are likely to be most useful for estimation.

First, we calibrate the risk free rate. The risk free rate is determined by many factors outside of the oil industry and consequently it would not be prudent to estimate it solely based on oil-price data. Also, it is clear by inspection that the risk-free rate is not identified by our model. It does not affect any of the moments we consider in our estimation and only appears in the constraint on model parameters in equation (A.7). Therefore, at best, futures price data can only impose a lower bound on the level of the risk-free rate, as implied by (A.7). Given all of the above considerations, we set the risk free rate at 2%.⁵

Next, consider a simple re-normalization of the structural parameters. Since futures prices in our model depend solely on the risk-neutral dynamics of the spot price, which in turn depends only on $\gamma\omega_t$, futures prices are determined by $\tilde{\omega}_t = \gamma\omega_t$, which evolves according to

$$d\tilde{\omega}_t = [-\gamma\mu^- + \gamma\bar{i}1_{\{\omega \leq \omega^*\}}] dt + \gamma\sigma_Y dW_t \quad (16)$$

where $1_{\{X\}}$ is the indicator function. Since we normalize the productivity parameter in (3) to one, only relative prices are informative, and therefore we can ignore the dependence of ω^* on structural parameters. Thus, the risk-neutral dynamics of futures prices is determined by only three combinations of five structural parameters: $\gamma\mu^-$, $\gamma\bar{i}$, and $\gamma\sigma_Y$. Therefore, we cannot identify all the model parameters separately from the futures data alone.

We obtain an additional identifying condition from the oil consumption data. As documented in Cooper (2003), world crude oil consumption increased by 46 per cent per capita from 1971 to 2000, implying an average growth rate of approximately 1.25%. Individual growth rates vary for the 23 three countries considered by Cooper, typically falling between -3 to 3% . For the US, the reported growth rate averaged -0.7% which we attempt to equate with the expected growth rate of oil consumption, g_C , implied by the model

$$g_C = \bar{i} \Pr(\omega \leq \omega^*) - \delta = \frac{1}{2}\sigma_Y^2 - (\lambda + \mu_Y), \quad (17)$$

where $\Pr(\omega \leq \omega^*) = \int_{-\infty}^{\omega^*} p^+(\omega) d\omega = (\bar{i})^{-1}\mu^-$ is the unconditional probability that ω is below the investment trigger.

⁵Our results regarding the 'V' shape response in prices are not affected by this choice.

Finally, to estimate the risk premium λ , we use average historic daily returns on fully collateralized futures positions (we use three-month contracts). We are thus left with five independent identifying restrictions on six structural parameters. Following Gomes (2001), we fix the depreciation rate of capital at $\delta = 0.12$ per year and do not infer it from futures prices. We estimate the remaining five parameters.

Parameter Estimates

Our estimated parameter values and the corresponding standard errors are summarized in Table 4. The first parameter value in the table, $\gamma = 3.92$, implies that the price elasticity of demand in our model is -0.26 . Cooper (2003) reports estimates of short-run and long-run demand elasticity for a partial adjustment demand equation based on US data of -0.06 and -0.45 respectively. In our model, there is no distinction between short-run and long-run demand, as demand adjustments are assumed to be instantaneous. Our estimate falls half-way between the two numbers reported in Cooper (2003) for the US and is close to the average of the long-run elasticity estimates reported for all 23 countries considered in that study, which is -0.2 .

Our second parameter is \bar{i} , the maximum investment rate in the model. This variable parameterizes the investment technology used by the firms. While it is difficult to make direct empirical comparisons, the upper bound of $\bar{i} = 0.19$ would allow for a plausible range of realized annual investment rates.

The average growth rate of demand is close to zero, as is the market price of risk. For comparison, the average annualized change in futures prices is approximately 2.8% in the data, which falls within the 95% confidence interval of the model's prediction. The volatility of demand shocks is not directly observable. The estimated value of σ_Y , together with the demand elasticity parameter γ^{-1} imply annualized volatility of the spot price of approximately 40%, which is closed to the observed price volatility of short-maturity futures contracts.

4.2 Results and Discussion

We first illustrate the fit of the model by plotting the term structure of unconditional futures price volatility (to facilitate comparison with empirical data, we express our results as daily values, defined as annual values scaled down by $\sqrt{252}$). We chose model parameters, as summarized in Table 4, to match the behavior of crude oil futures. Figure 7 compares the volatility of prices implied by our choice of parameters to the empirical estimates. Our model seems capable of reproducing the slow-decaying pattern of futures price volatility. This feature of the data presents a challenge to simple storage models, as discussed in Routledge et. al. (2000)). To see why it may not be easy to reproduce the slow-decaying pattern of unconditional volatilities in a simple single-factor model, consider a reduced-form model in which the logarithm of the spot price process follows a continuous-time AR(1) process (Ornstein-Uhlenbeck process). Specifically, assume that the spot price is given by

$$S_t = e^{y_t}. \quad (18)$$

and under the risk-neutral probability measure y_t follows

$$dy_t = \theta_y(\bar{y} - y_t)dt + \sigma_y dW_t, \quad (19)$$

where θ_y is the mean-reversion coefficient and \bar{y} is the long-run mean of the state variable. According to this simple model, the unconditional volatility of futures price changes is an exponential function of maturity τ :

$$\sigma^2(\tau) = \sigma_y^2 e^{-2\theta\tau}. \quad (20)$$

To compare the term structure of unconditional volatility implied by this model to the one generate by our model, we calibrate parameters θ_y and σ_y so that the simple model exhibits the same volatility of the spot price and the same 30-day autocorrelation of the basis as our model. Figure 7 shows that, as expected, unconditional volatility implied by the simple model above decays too fast relative to our model and data.

The main qualitative distinction between the properties of our model and those of basic storage models is in the conditional behavior of futures volatility. As we demonstrate in Section 2, the empirical relation between the volatility of futures prices and the slope of the

term structure of prices is non-monotone and has a pronounced “V-shape”. Intuitively, we would expect our model to exhibit this pattern. When the state variable ω_t is far away from the investment trigger ω^* , one of the investment constraints is binding and can be expected to remain binding for some time. If the capital stock K_t is much higher than its optimal level, given the current level of demand, firms find it optimal to postpone investment and the irreversibility constraint binds. On the other hand, when K_t is much lower than the optimal level, firms invest at the maximum possible rate and the investment rate constraint binds. In either case, the supply of the commodity is relatively inelastic and futures prices are relatively volatile. The further ω_t travels away from the investment trigger, the larger the effect on volatility of long-maturity futures. At the same time, it is precisely when ω_t is relatively far away from the investment trigger ω^* , when the absolute value of the slope of the term structure of futures prices is large. This is to be expected. All prices in our model are driven by a single mean-reverting stationary state variable, and since futures prices of longer-maturity contracts are less sensitive to the current value of the state variable than the spot price, the slope of the forward curve tends to be large when the state variable is far away from its long-run average value. The latter, in turn, is not far from ω^* , given that ω_t reverts to ω^* . Thus, our model predicts that the volatility of futures prices should exhibit a “V-shape” as a function of the slope of the term structure of futures prices.

It should be clear from the above discussion that the critical feature of the model is not the precise definition of the production function, but rather the variable-elasticity property of the supply side of the economy. The “V-shape” pattern in volatilities is due to the fact that supply can adjust relatively easily in response to demand shocks when the output is close to the optimal level, but supply is relatively inelastic when the output level is far from the optimum.

We now report the quantitative properties of the model. The long-run average of the slope of the forward curve, $\ln \left[\frac{P(t-1,3)}{P(t-1,1)} \right]$, is 0.01 in the model, compared to the empirical value of -0.0151 . Both values are statistically indistinguishable from zero. The long-run standard deviation of the slope in the model, which equals 0.0287, is slightly below the empirical value of 0.0326. The 30-day autocorrelation coefficient of the slope implied by the model is equal to 0.80, as compared to the value of 0.74 in the data. Overall, our

model fits the basic behavior of the slope of the forward curve quite well. Table 5 shows the estimates of linear and piece-wise linear specifications of conditional variance of futures price changes (2) implied by the model for one-, five-, and ten-month futures. The coefficients of linear regressions are negative and close in magnitude to their empirical counterparts. Such a negative relation between conditional volatility of futures prices and the basis would typically be interpreted as supportive of simple storage models. Note, however, that our model without storage can reproduce the same kind of relation. Our model, however, has a further important implication: the linear model is badly misspecified, since the theoretically predicted relation is non-monotone. Our piece-wise linear specification produces coefficients $\beta_{1,T}$ and $\beta_{2,T}$ that agree with their empirical counterparts in sign but differ in magnitude. Given the extremely streamlined nature of our model (e.g., the basis is a sufficient statistic for conditional volatility), this should not be surprising. Also, it is important to keep in mind that we did not target the volatility-basis relation in our estimation of structural parameters. The entire distribution of regression coefficients across maturities of the futures contracts is shown in Figure 8. Finally, Figure 9 helps visualize the “V-shape” pattern.

In order to understand the sensitivity of our results to the baseline parameters summarized in Table 4, we compute elasticities of basic statistics of the model output with respect to these parameters. Each elasticity is calculated by simulating the model twice: with a value of the parameter of interest one standard deviation below (above) its baseline value. Next, the change in the moment is calculated as the difference between the results from the two simulations. This difference is then divided by the change in the underlying structural parameter between the two simulations. Finally, the result is then multiplied by the ratio of the baseline structural parameter to the baseline moment. The elasticities are reported in Table 6.

An increase in the demand volatility, σ_Y , or in the elasticity of the inverse demand curve, γ , leads to an increase in the volatility of the spot price, which equals $\gamma^2 \sigma_Y^2$. As one would expect, volatility of futures prices of various maturities increases as well. Qualitatively, both of the parameters σ_Y and γ affect the level of the unconditional volatility curve plotted in Fig. 7. However, the demand volatility has strong positive effect on the expected growth rate of oil consumption since it increases the long-run growth rate of the level of the demand

curve, Y_t^{-1} . γ has no such effect.

The constraint on the investment rate \bar{i} has no effect on the volatility of the spot price. However, it affects volatility of futures prices. A higher value of \bar{i} allows capital stock to adjust more rapidly in response to positive demand shocks, thus reducing the impact of demand shocks on the future value of the spot price and therefore lowering the volatility of futures prices. We thus see that \bar{i} effectively controls the slope of the term structure of volatility, higher values of \bar{i} imply a steeper term structure. \bar{i} has no effect on the expected growth rate of oil consumption, in agreement with Eq. (17).

An increase in the unconditional mean of the demand shock, μ_Y , has little affect on the level of futures price volatility. This is not surprising given the role μ_Y plays in the evolution of the state variable ω_t . An increase in μ_Y raises the drift of ω_t uniformly. The impact of this on the volatility of futures price is ambiguous and depends on the relative magnitude of the drift of ω_t , μ_ω , above and below the investment threshold ω^* . By symmetry considerations, if $\mu^+ = \mu^-$, an infinitesimal change in μ_Y has no impact on the volatility of futures prices. Under the calibrated parameter values, $\mu^- = 0.11$ and $\mu^+ = 0.08$ and futures volatility is not very sensitive to μ_Y . The same is true for the risk premium, λ . Both μ_Y and λ have strong negative effect on g_C in agreement with Eq. (17).

In general, the affect of model parameters on the slope of the forward curve is difficult to interpret intuitively and depends on the chosen parameter values. However, the fact that the moments of the slope have different sensitivities to various model parameters makes them useful in estimating these parameters.

Finally, as a robustness check of our model we consider a possibility that a “V-shape” in conditional volatility could arise mechanically due to heteroscedasticity in demand. We therefore construct an alternative theoretical specification in which shocks to demand are heteroscedastic, and their instantaneous volatility follows a continuous-time Markov process, independent of other exogenous processes in the economy. We also assume that supply is constant, while the level of demand is a mean-reverting stationary process. We find that not to be the case – independent shocks to demand volatility cannot give rise to a non-monotone relation between the slope of the forward curve and conditional volatility of futures prices. Details of this analysis are presented in Appendix B.

5 Conclusions

This paper contributes along two dimensions. First, we show that volatility of future prices has a “V-shape” relationship with respect to the slope of the term structure of futures prices. Second, we show that such volatility patterns arise naturally in models that emphasize investment constraints and, consequently, time-varying supply-elasticity as a key mechanism for price dynamics. Our empirical findings seem beyond the scope of simple storage models, which are currently the main focus of the literature, and point towards investigating alternative economic mechanisms, such as the one analyzed in this paper. Future work will entail a model that nests both storage and investment in an attempt to isolate their quantitative effects.

References

- Brennan, M., 1991, “The Price of Convenience and the Valuation of Commodity Contingent Claims” in D. Lund and B. Oksendal (eds.), *Stochastic Models and Option Models*, Elsevier Science Publishers, 1991.
- Brennan, M. and E. Schwartz, 1985, “Evaluating Natural Resource Investments,” *Journal of Business* 58, 135–157.
- Casassus, J., P. Collin-Dufresne and B. Routledge, 2004, “Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technologies”, CMU working paper.
- Cooper, John C. B., 2003, “Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries,” *OPEC Review* 27, 1-8.
- Deaton, A. and G. Laroque, 1992, “On the Behavior of Commodity Prices,” *The Review of Financial Studies* 59, 1–23.
- Deaton, A. and G. Laroque, 1996, “Competitive Storage and Commodity Price Dynamics,” *Journal of Political Economy* 104, 896–923.
- Duffie, Darrell and Ken Singleton, 1993, “Simulated moments estimation of Markov models of asset prices,” *Econometrica* 61, 929-952.
- Gomes, Joao F., 2001, “Financing investment,” *American Economic Review* 91, 1263-1285.
- Kaldor, N., 1939, “Speculation and Economic Stability,” *The Review of Economic Studies* 7, 1–27.
- Kogan, L., 1999, Doctoral Dissertation, Sloan School of Management, MIT, Boston, MA.
- Lee, B. and B. Ingram, 1991, “Simulation estimation of time series models”, *Journal of Econometrics* 47, 197-205.
- Lucas, R., and E. Prescott, 1971, “Investment under Uncertainty,” *Econometrica* 39, 659–681.

- Routledge, B., D. Seppi and C. Spatt, 2000, "Equilibrium Forward Curves for Commodities," *Journal of Finance* 55, 1297–1338.
- Schwartz, E., 1997, "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging," *Journal of Finance* 52, 923–973.
- Telser, L.G., 1958, "Futures Trading and the Storage of Cotton and Wheat," *Journal of Political Economy* 66, 133–144.
- Williams, J. and B. Wright, 1991, *Storage and Commodity Markets*. Cambridge, England: Cambridge University Press.
- Working, H., 1948, "Theory of the Inverse Carrying Charge in Futures Markets," *Journal of Farm Economics* 30, 1–28.
- Working, H., 1949, "The Theory of the Price of Storage," *American Economic Review* 39, 1254–1262.

Appendix

A. Central Planner's Problem

To characterize the solution of the social planner's problem, we define the value function $J(K_t, Z_t)$ as

$$J(K_t, Z_t) = \max_{i_s \in [0, \bar{i}]} E_t \left[\int_t^\infty e^{-r_f(s-t)} \left(Y_s^{-\gamma} \frac{K_s^{1-\gamma}}{1-\gamma} - i_s K_s \right) ds \right].$$

The Hamilton-Jacobi-Bellman (HJB) equation for $J(K_t, Z_t)$ takes the form

$$\max_{i_t \in [0, \bar{i}]} [i K J_K - i K] + \frac{\sigma_Y^2}{2} Y^2 J_{YY} + \mu_Y Y J_Y - \delta K J_K + Y^{-\gamma} \frac{K^{1-\gamma}}{1-\gamma} - r_f J = 0. \quad (\text{A.1})$$

The solution to the above equation consists of the value function and endogenously determined no-investment boundary determined by $J_K = 1$. When $J_K \geq 1$ the investment is made at a maximum rate $i = \bar{i}$, and no investment is made when $J_K < 1$.

We make the following change of variables in HJB equation

$$k = \ln K, \quad y = \ln Y,$$

to obtain

$$\max_{i_t \in [0, \bar{i}]} [i J_k - i e^{-y} e^{k+y}] + \frac{\sigma_Y^2}{2} J_{yy} + (\mu_Y - \frac{\sigma_Y^2}{2}) J_y - \delta J_k + \frac{e^{-y} e^{(1-\gamma)(k+y)}}{1-\gamma} - r_f J = 0.$$

We will look for the solution in the form

$$J(k, y) = \frac{e^{-y} f(\omega)}{1-\gamma}, \quad (\text{A.2})$$

where $\omega = k + y$. We have

$$\begin{aligned} J_k &= \frac{e^{-y} f'}{1-\gamma}, \\ J_y &= \frac{e^{-y}}{1-\gamma} (f' - f), \\ J_{yy} &= \frac{e^{-y}}{1-\gamma} (f'' - 2f' + f), \end{aligned}$$

and the ODE for f takes the form

$$\max_{i \in [0, \bar{i}]} [i f' - i (1-\gamma) e^\omega] + \frac{\sigma_Y^2}{2} f'' + (\mu_Y - \delta - \frac{3\sigma_Y^2}{2}) f' + (\sigma_Y^2 - \mu_Y - r_f) f + e^{(1-\gamma)\omega} = 0. \quad (\text{A.3})$$

We will look for a solution of (A.1) of the form

$$f(\omega) = \begin{cases} A \exp(\kappa_-(\omega - \omega^*)) + B \exp((1 - \gamma)\omega), & \omega \geq \omega^*, \\ \bar{A} \exp(\bar{\kappa}_+(\omega - \omega^*)) + \bar{B} \exp((1 - \gamma)\omega) + C \exp(\omega), & \omega < \omega^*, \end{cases}, \quad (\text{A.4})$$

for all γ 's, where ω^* is the no-investment boundary in terms of the new state variable ω .

Substituting these solutions into ODE (A.3) yields quadratic equations on κ_{\pm} and $\bar{\kappa}_{\pm}$

$$\begin{aligned} \frac{\sigma_Y^2}{2} \kappa_{\pm}^2 - \left(\delta + \frac{3\sigma_Y^2}{2} - \mu_Y\right) \kappa_{\pm} + (\sigma_Y^2 - \mu_Y - r_f) &= 0, \\ \frac{\sigma_Y^2}{2} \bar{\kappa}_{\pm}^2 - \left(\delta - \bar{i} + \frac{3\sigma_Y^2}{2} - \mu_Y\right) \bar{\kappa}_{\pm} + (\sigma_Y^2 - \mu_Y - r_f) &= 0, \end{aligned} \quad (\text{A.5})$$

and constants B , \bar{B} , and C

$$\begin{aligned} B &= \frac{1}{r_f + \gamma\mu_Y + (1 - \gamma)\delta - \frac{\gamma(1+\gamma)}{2}\sigma_Y^2}, \\ \bar{B} &= \frac{1}{r_f + \gamma\mu_Y - (1 - \gamma)(\bar{i} - \delta) - \frac{\gamma(1+\gamma)}{2}\sigma_Y^2}, \\ C &= \frac{(1 - \gamma)\bar{i}}{\bar{i} - \delta}. \end{aligned} \quad (\text{A.6})$$

For the quadratic equation on κ to have one negative root, we impose that

$$r_f + \mu_Y - \sigma_Y^2 > 0. \quad (\text{A.7})$$

As of now we have determined constants B , \bar{B} , C , κ_- and $\bar{\kappa}_+$, and we are left to find four more parameters, constants A , \bar{A} , and, finally, the no-investment boundary ω^* in both cases. We have three boundary conditions to determine these constants. The first two boundary conditions ensure continuity of $f(\omega)$ and its first derivative across the no-investment boundary

$$\begin{aligned} f(\omega^* - 0) &= f(\omega^* + 0), \\ f'(\omega^* - 0) &= f'(\omega^* + 0). \end{aligned} \quad (\text{A.8})$$

The investment optimality condition, $J_K = 1$, which takes the following form in terms ω

$$f'(\omega^*) = (1 - \gamma)e^{\omega^*}. \quad (\text{A.9})$$

is the third and the final one. Substituting equation (A.4) into the boundary conditions we obtain the following system of equations

$$\begin{cases} \bar{A} - A = (B - \bar{B}) e^{(1-\gamma)\omega^*} - C e^{\omega^*} \\ \bar{\kappa}_+ \bar{A} - \kappa_- A = (1 - \gamma) (B - \bar{B}) e^{(1-\gamma)\omega^*} - C e^{\omega^*} \\ \kappa_- A + (1 - \gamma) B e^{(1-\gamma)\omega^*} = (1 - \gamma) e^{\omega^*} \end{cases}, \quad (\text{A.10})$$

Solving last two equations in both systems we find the unique solution

$$A = \frac{\bar{\kappa}_+ + \gamma - 1}{\bar{\kappa}_+ + |\kappa_-|} (\bar{B} - B) e^{(1-\gamma)\omega^*} + \frac{\bar{\kappa}_+ - 1}{\bar{\kappa}_+ + |\kappa_-|} C e^{\omega^*}, \quad (\text{A.11})$$

$$\begin{aligned} \bar{A} &= \frac{\gamma - |\kappa_-| - 1}{\bar{\kappa}_+ + |\kappa_-|} (\bar{B} - B) e^{(1-\gamma)\omega^*} - \frac{|\kappa_-| + 1}{\bar{\kappa}_+ + |\kappa_-|} C e^{\omega^*}, \quad (\text{A.12}) \\ e^{\gamma\omega^*} &= \frac{(1 - \gamma + |\kappa_-|) \bar{\kappa}_+ B + (1 - \bar{\kappa}_+) |\kappa_-| \bar{B}}{(1 - \gamma) (\bar{\kappa}_+ + |\kappa_-|) + |\kappa_-| (\bar{\kappa}_+ - 1) C}. \end{aligned}$$

B. Stationary long-run distribution of ω_t

In this appendix we calculate the stationary long-run distribution of the state variable ω , dynamics of which can be written as

$$d\omega_t = [\mu^+ 1_{\{\omega \leq \omega^*\}} - \mu^- 1_{\{\omega > \omega^*\}}] dt + \sigma_Y dW_t. \quad (\text{B.1})$$

The stationary distribution of ω

$$p(\omega) = \begin{cases} p^+(\omega) & \omega \leq \omega^* \\ p^-(\omega) & \omega > \omega^* \end{cases}, \quad (\text{B.2})$$

exists if $0 < \mu^- < \bar{i}$ and it satisfies backward Kolmogorov ODE

$$\frac{d^2 p(\omega)}{d\omega^2} - \frac{2 [\mu^+ 1_{\{\omega \leq \omega^*\}} - \mu^- 1_{\{\omega > \omega^*\}}]}{\sigma_Y^2} \frac{dp(\omega)}{d\omega} = 0. \quad (\text{B.3})$$

$p(\omega)$ also satisfies the normalization condition

$$\int_{-\infty}^{\omega^*} p^+(\omega) d\omega + \int_{\omega^*}^{\infty} p^-(\omega) d\omega = 1. \quad (\text{B.4})$$

Condition (B.4) eliminates a constant as a trivial solution of the ODE (B.3) which can be integrated once to yield

$$\frac{dp(\omega)}{d\omega} - \frac{2 [\mu^+ 1_{\{\omega \leq \omega^*\}} - \mu^- 1_{\{\omega > \omega^*\}}]}{\sigma_Y^2} p(\omega) = 0. \quad (\text{B.5})$$

We can find necessary boundary condition at $\omega = \omega^*$ by integrating equation (B.5) over the interval $\omega \in [\omega^* - \varepsilon, \omega^* + \varepsilon]$ and then taking limit of $\varepsilon \rightarrow 0$

$$\lim_{\varepsilon \rightarrow 0} \int_{\omega^* - \varepsilon}^{\omega^* + \varepsilon} \left[\frac{dp(\omega)}{d\omega} - \frac{2 [\mu^+ 1_{\{\omega \leq \omega^*\}} - \mu^- 1_{\{\omega > \omega^*\}}]}{\sigma_Y^2} p(\omega) \right] d\omega = p^-(\omega^*) - p^+(\omega^*) = 0. \quad (\text{B.6})$$

Therefore, we obtain that $p(\omega)$ is continuous at $\omega = \omega^*$. We can now solve ODE (B.3) separately for $p^+(\omega)$ and $p^-(\omega)$

$$\begin{aligned} p^+(\omega) &= A^+ e^{\frac{2\mu^+}{\sigma_Y^2}\omega}, \\ p^-(\omega) &= A^- e^{-\frac{2\mu^-}{\sigma_Y^2}\omega}, \end{aligned} \tag{B.7}$$

where A^\pm can be found from the boundary and normalization conditions which take the following forms

$$\begin{aligned} A^+ e^{\frac{2\mu^+}{\sigma_Y^2}\omega^*} &= A^- e^{-\frac{2\mu^-}{\sigma_Y^2}\omega^*}, \\ \frac{\sigma_Y^2}{2\mu^+} A^+ e^{\frac{2\mu^+}{\sigma_Y^2}\omega^*} + \frac{\sigma_Y^2}{2\mu^-} A^- e^{-\frac{2\mu^-}{\sigma_Y^2}\omega^*} &= 1. \end{aligned} \tag{B.8}$$

Solving equations (B.8) we finally obtain Eq. (11).

C. Robustness Check

In this appendix we analyze a simple economy with heteroscedastic shocks to the level of the demand curve. Our objective is to establish that the ‘‘V-shape’’ of the conditional volatility of futures prices implied by our production-based model does not arise spuriously simply due to heteroscedasticity in demand shocks. An informal argument would suggest that such a possibility exists. If demand shocks have time-varying volatility, one could expect the slope of the forward curve to move further away from its long-run mean when conditional volatility of demand shocks is relatively high, thus giving rise to a ‘V-shaped’ relation observed in the data. The above argument is informal because it makes an implicit connection between the volatility of demand shocks and the volatility of the slope of the forward curve, meanwhile assuming that the conditional mean of the slope is unaffected by changes in demand shock volatility. Below we develop a simple explicit model and show that heteroscedasticity of demand shocks does not immediately lead to a convex relation between the conditional futures price volatility and the slope of the demand curve.

Our model here is similar to the main model in the paper. As before, consumers represented by the demand curve

$$Q_t = Y_t^{-1} S_t^{-\frac{1}{\gamma}}, \tag{C.1}$$

where unexpected changes in Y_t represent demand shocks. The most significant departure from the earlier formulation is the simplifying assumption that the level of supply is fixed at $Q = e^q$. To make sure that prices are stationary, we assume that the level of demand curve $y_t = \ln Y_t$ follows a linear mean-reverting process

$$dy_t = \theta(\bar{y} - y_t)dt + \sigma_t dW_t, \quad (\text{C.2})$$

The conditional volatility of the demand process is stochastic, it follows a continuous Markov process independent of W_t . The precise functional form of the volatility process is not important for our results.

In equilibrium, supply equals demand and hence the spot price is equal to

$$S_t = e^{-\gamma(y_t + q)}, \quad (\text{C.3})$$

while futures price is equal to

$$P(t, T) = E_t[S_T] = e^{-\gamma q} E_t[e^{-\gamma y_T}], \quad \forall T \geq 0. \quad (\text{C.4})$$

We now characterize futures prices analytically. Notice that since the process for volatility is independent of demand shocks, we can first calculate the expectation over W_t , treating the process σ_t as deterministic, and then calculate the expectation over σ_t .

The solution for y_T is standard:

$$y_T = y_t Z(t) + \bar{y}(1 - Z(t)) + \int_t^T Z(s) \sigma_s dW_s, \quad (\text{C.5})$$

where $Z(t) \equiv e^{-\theta(T-t)}$. Taking expectations over W_t we find that futures prices are given by

$$P(t, T) = S_t^{Z(t)} e^{-\gamma(\bar{y} + q)(1 - Z(t))} E \left[\exp \left\{ \frac{\gamma^2}{2} \int_t^T Z^2(s) \sigma_s^2 ds \right\} \middle| \sigma_t \right]. \quad (\text{C.6})$$

It is easy to see that the conditional volatility of futures prices of all maturities is equal to σ_t . The slope of the forward curve is given by

$$\frac{P(t, T)}{S_t} = S_t^{Z(t)-1} e^{-\gamma(\bar{y} + q)(1 - Z(t))} E \left[\exp \left\{ \frac{\gamma^2}{2} \int_t^T Z^2(s) \sigma_s^2 ds \right\} \middle| \sigma_t \right]. \quad (\text{C.7})$$

Because the volatility process σ_t is continuous and Markov, the conditional expectation in the previous expression is increasing in σ_t . Thus, we find that the conditional volatility of futures prices is positively related to slope of the term structure, in contrast to a non-monotone convex relation implied by the main model in the paper.

Figure 1: **Term structure of volatility, crude oil futures**

The data are daily percentage price changes on NYMEX crude oil (CL) futures from 1983 to 2000. σ denotes the standard deviation of daily percentage price changes. The time to maturity is defined as number of months left until the delivery month. The unconditional standard deviation is constructed using sample's first and second moments, while standard deviations conditional on backwardation and contango are conditioned on the shape of the forward curve one day prior.

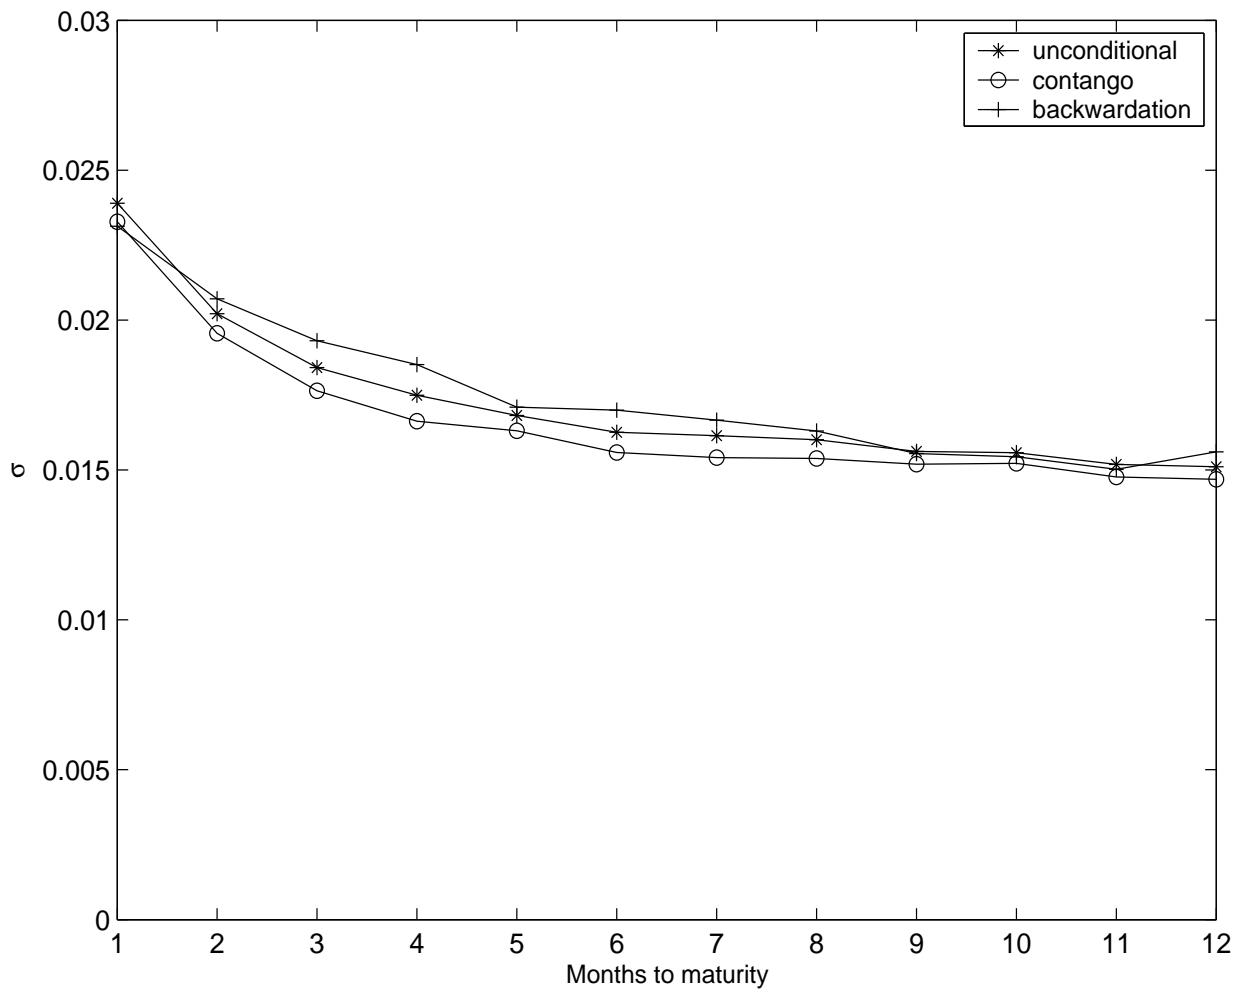


Figure 2: **Term structure of volatility, heating oil futures**

The data are daily percentage price changes on NYMEX heating oil (HO) futures from 1979 to 2000. See the caption to Figure 1 for details.

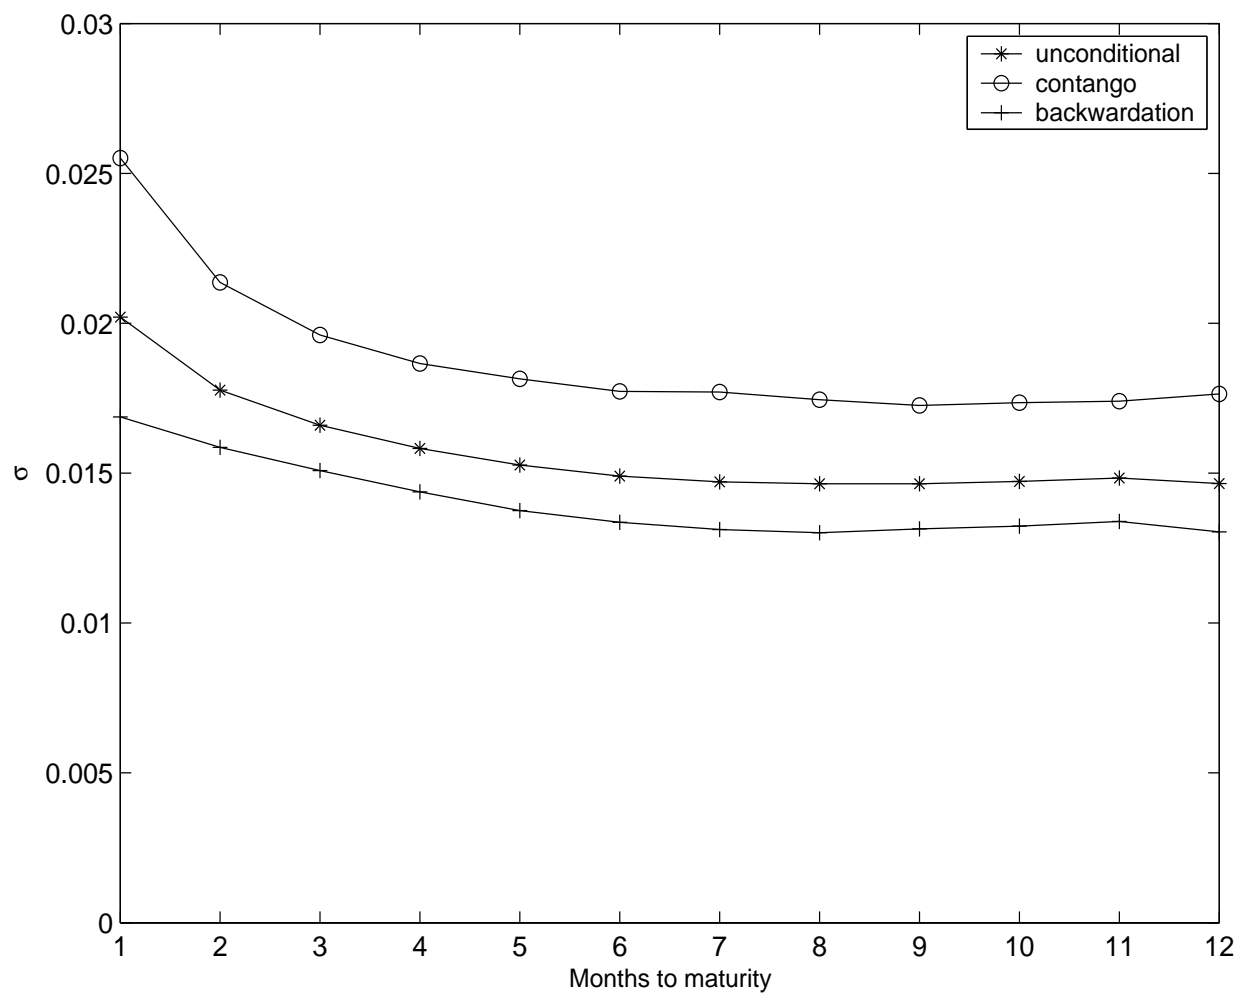


Figure 3: **Term structure of volatility, unleaded gasoline futures**

The data are daily percentage price changes on NYMEX unleaded gasoline (HU) futures from 1985 to 2000. See the caption to Figure 1 for details.

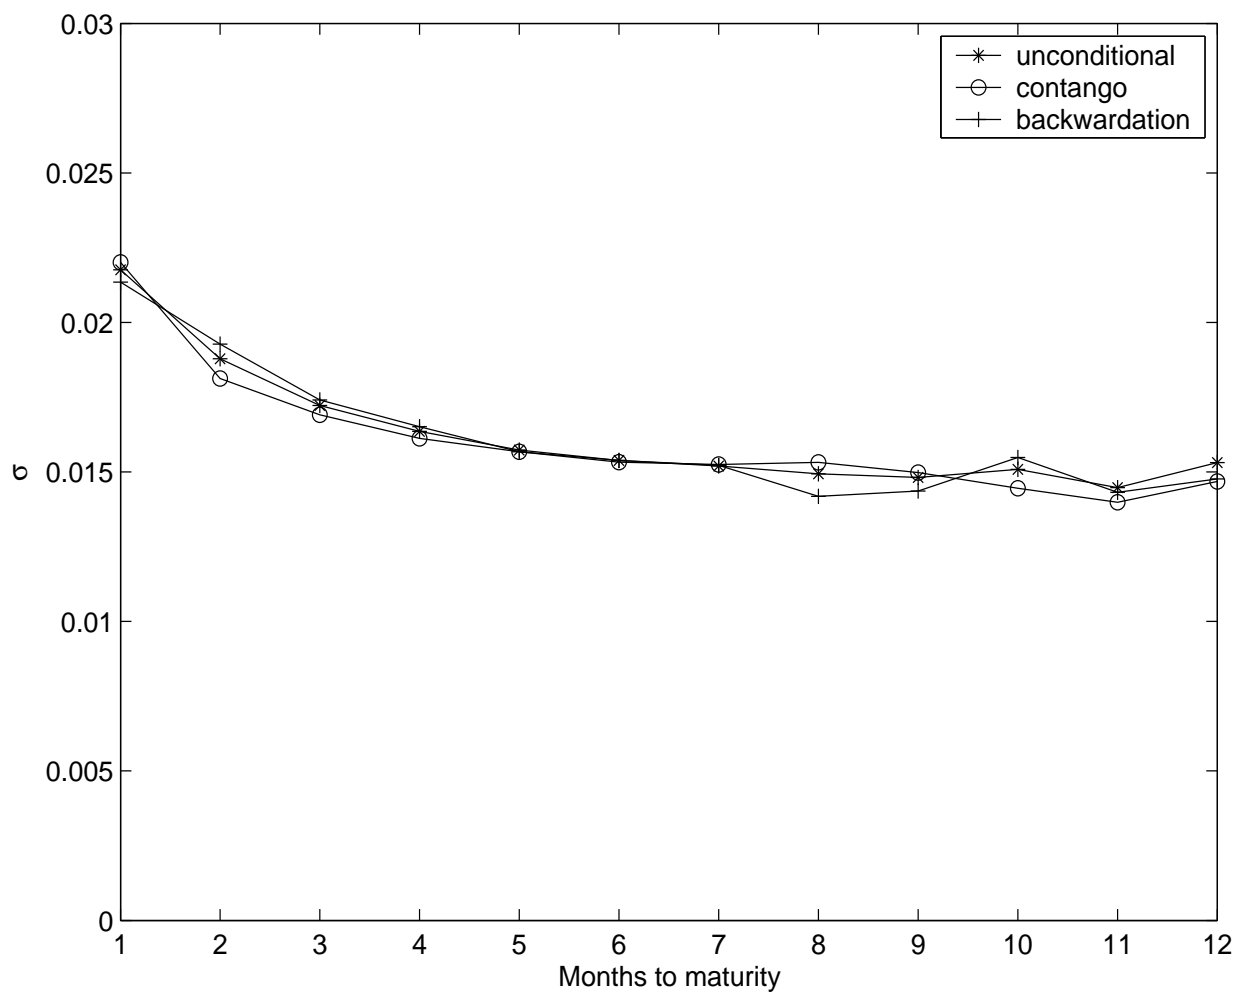


Figure 4: **Conditional volatility, crude oil futures**

The data are daily percentage price changes on NYMEX crude oil (CL) futures from 1983 to 2000, denoted by $R(t, T)$. Two different specifications are used to relate the instantaneous volatility of futures prices to the beginning-of-period slope of the forward curve. The first specification is

$$|R(t, T)| = \alpha_T + \beta_T \ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] + \varepsilon_T(t).$$

The second specification decomposes the slope into positive and negative parts

$$|R(t, T)| = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] \right)^- + \varepsilon_T(t),$$

where $(X)^\pm$ denotes positive (negative) part of X respectively. The figure shows the estimates of all three coefficients (exhibit A) and their respective t -statistics (exhibit B) for different times to maturity. All t -statistics are White-adjusted.

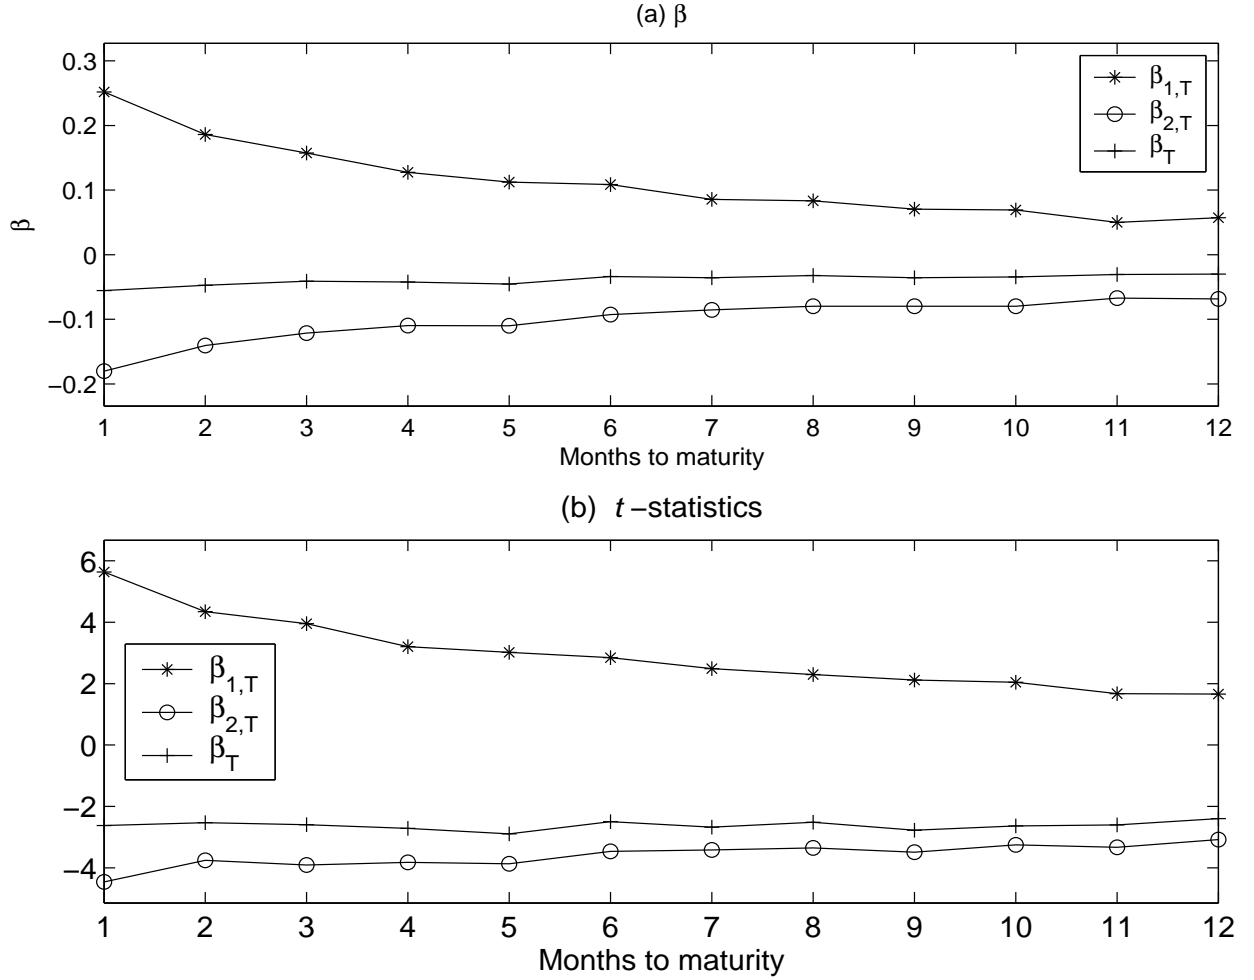


Figure 5: **Conditional volatility, heating oil futures**

The data are daily percentage price changes on NYMEX heating oil (HO) futures from 1979 to 2000. See the caption to Figure 4 for a detailed description.

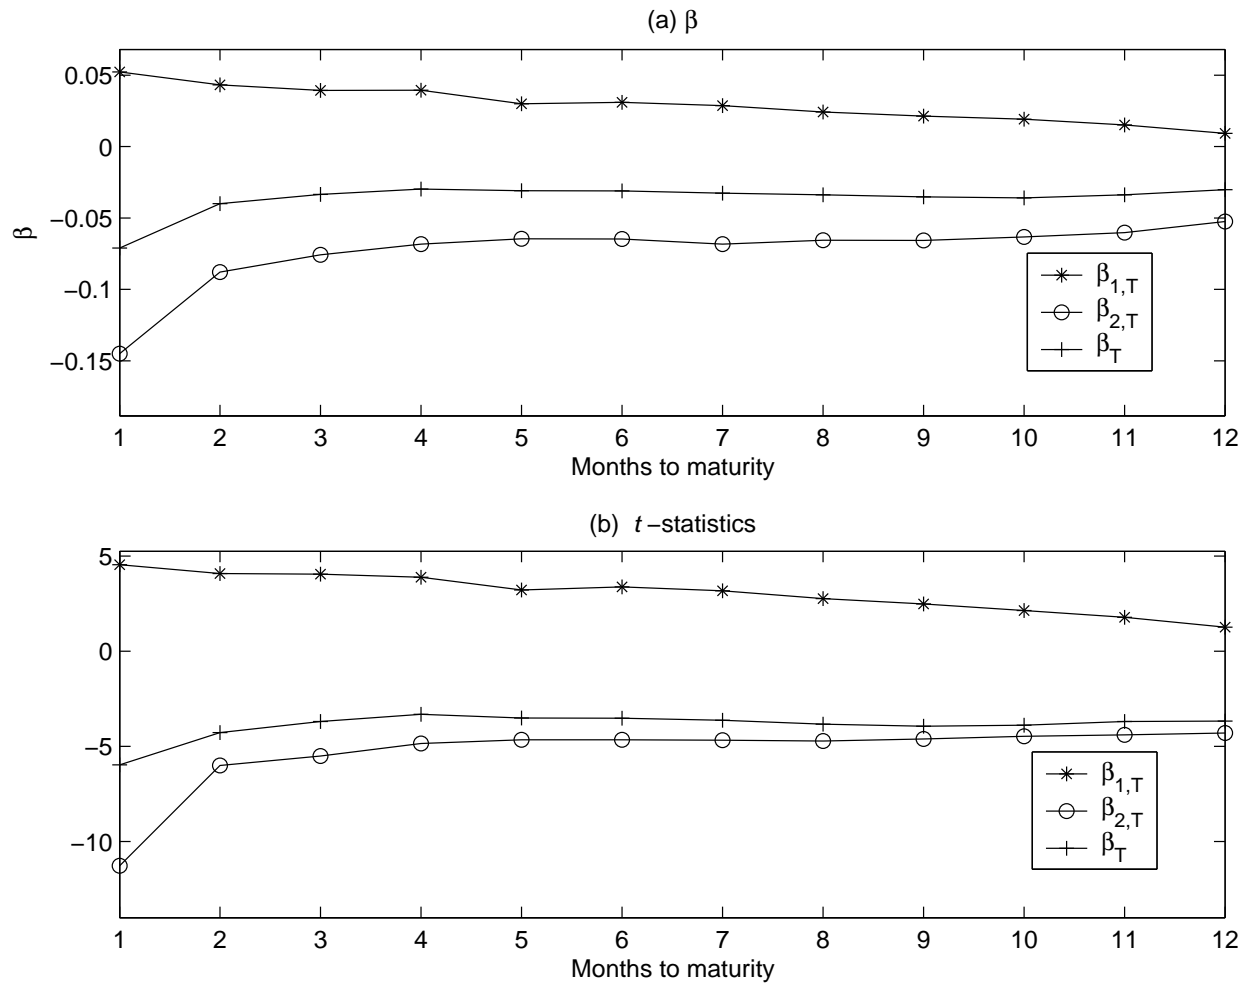


Figure 6: **Conditional volatility, unleaded gasoline futures**

The data are daily percentage price changes on NYMEX unleaded gasoline (HU) futures from 1985 to 2000. See the caption to Figure 4 for a detailed description.

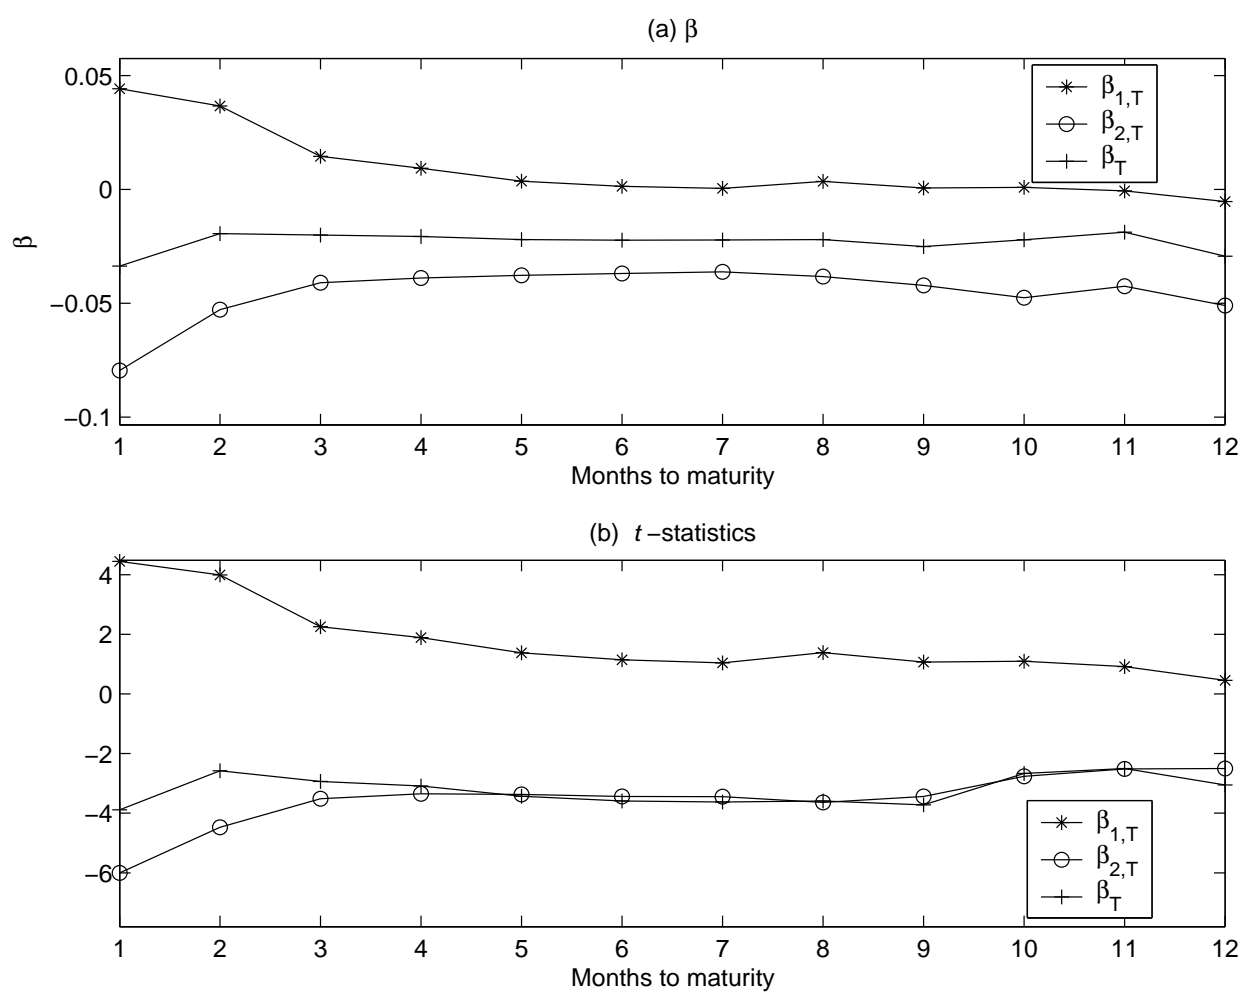


Figure 7: **Unconditional volatility of futures prices, model output**

The unconditional standard deviation of daily changes in futures prices is constructed based on the output from the model (model 1). The model is fitted to the data on crude oil futures (CL) using a two-step simulated method of moments described in Section 4. Parameter values are reported in Table 4. Model 2 corresponds to the unconditional standard deviation of changes in futures prices implied by Eq. (20).

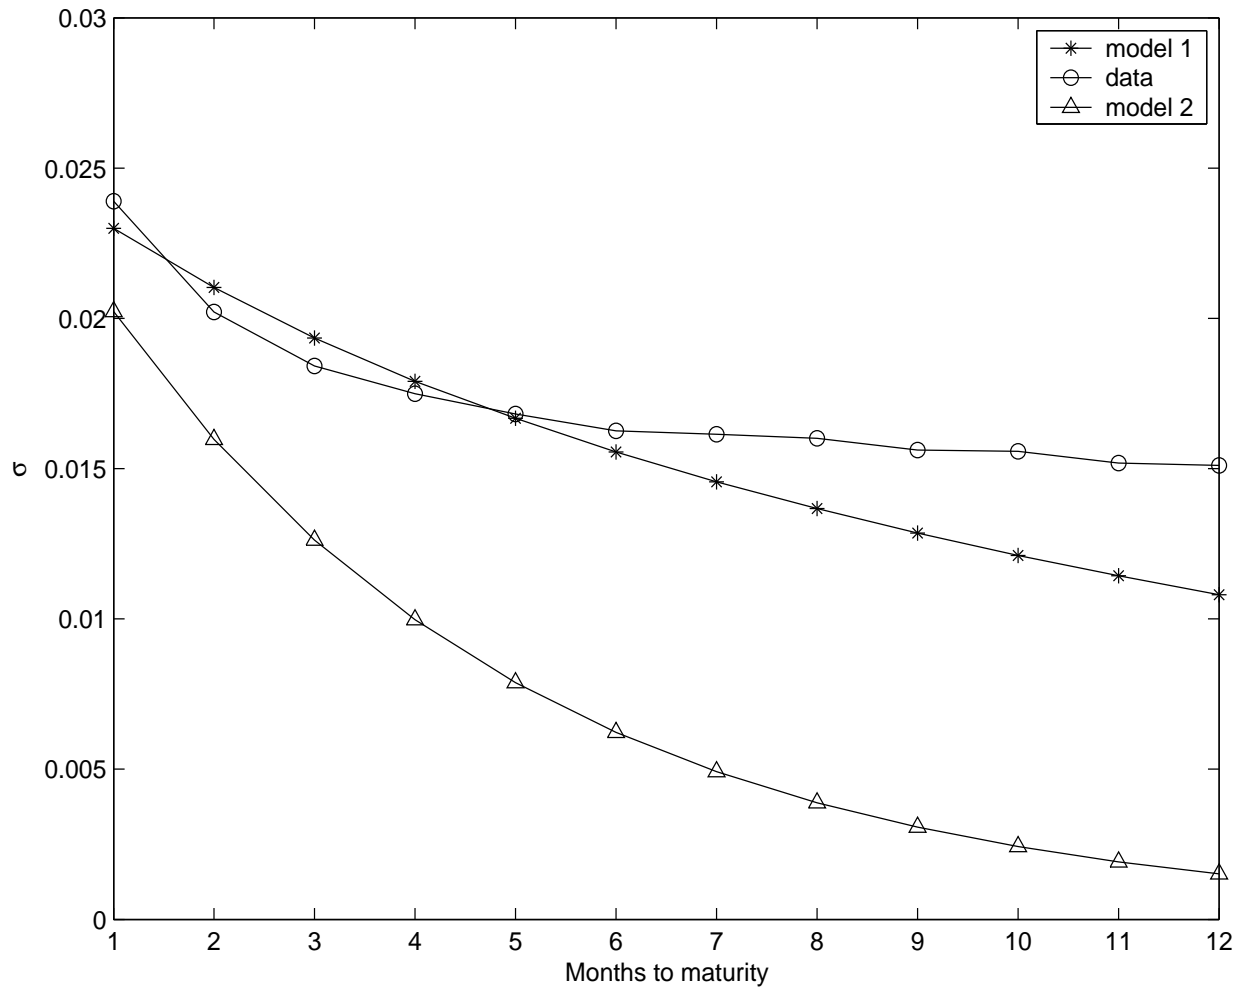


Figure 8: **Conditional volatility of futures prices, model output**

For each time to maturity, T , we simulate a time series of daily futures prices using parameter values reported in Table 4. We then compute daily percentage changes in futures prices, denoted by $R(t, T)$. Two different specifications are used to relate the instantaneous volatility of futures prices to the beginning-of-period slope of the forward curve. The first specification is

$$|R(t, T)| = \alpha_T + \beta_T \ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] + \varepsilon_T(t).$$

The second specification decomposes the slope into positive and negative parts

$$|R(t, T)| = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(t-1, 3)}{P(t-1, 1)} \right] \right)^- + \varepsilon_T(t),$$

where $(X)^\pm$ denotes positive (negative) part of X respectively. The figure shows all three betas for different times to maturity. See the caption to Figure 4 for further details.

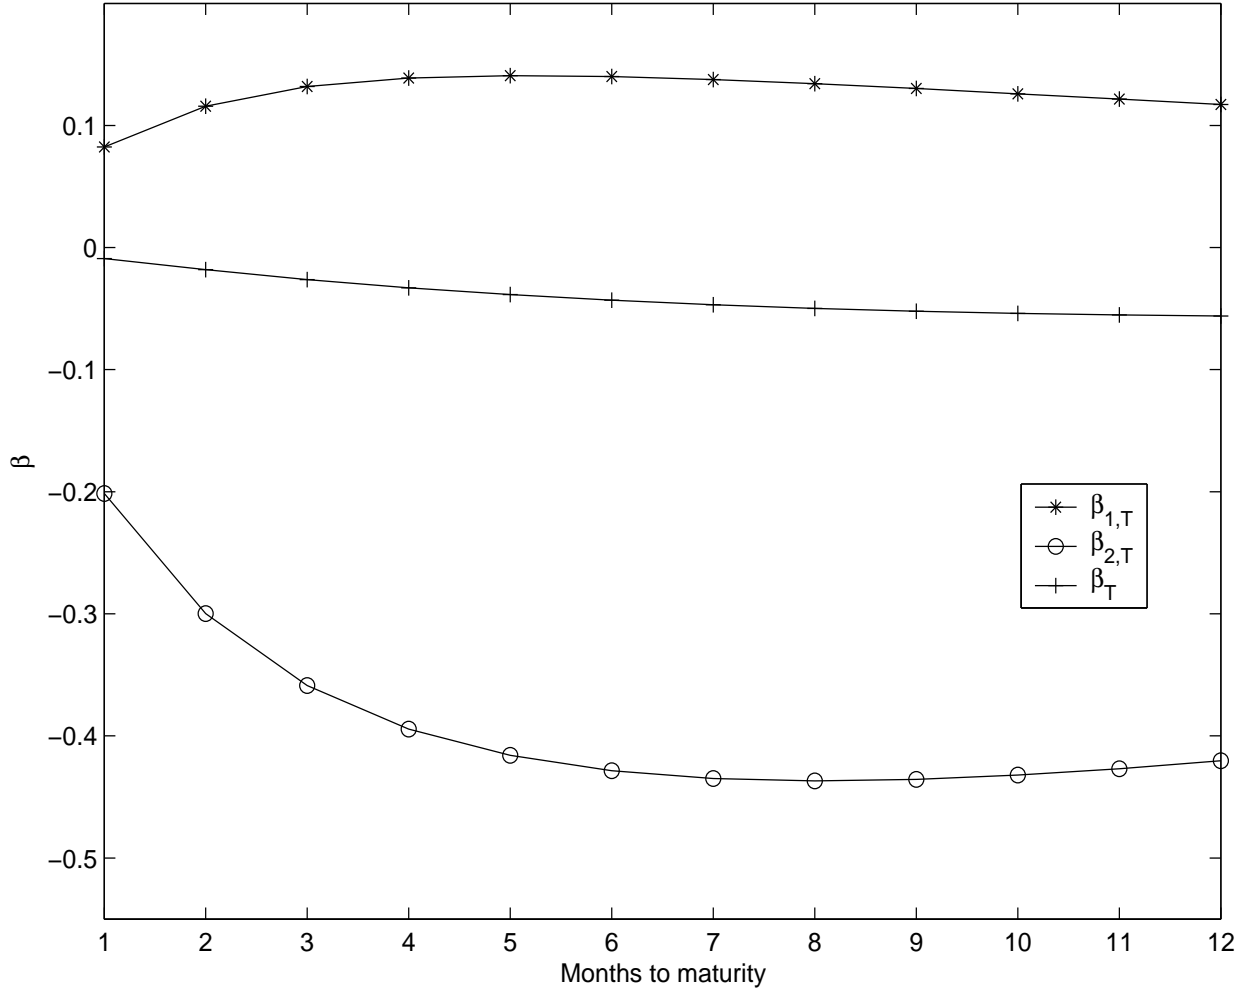


Figure 9: **Properties of volatility of futures prices, model output**

The figure reports the model's term structure of futures volatility with respect to time of maturity and the slope of the forward curve. Volatility is expressed in annual terms.

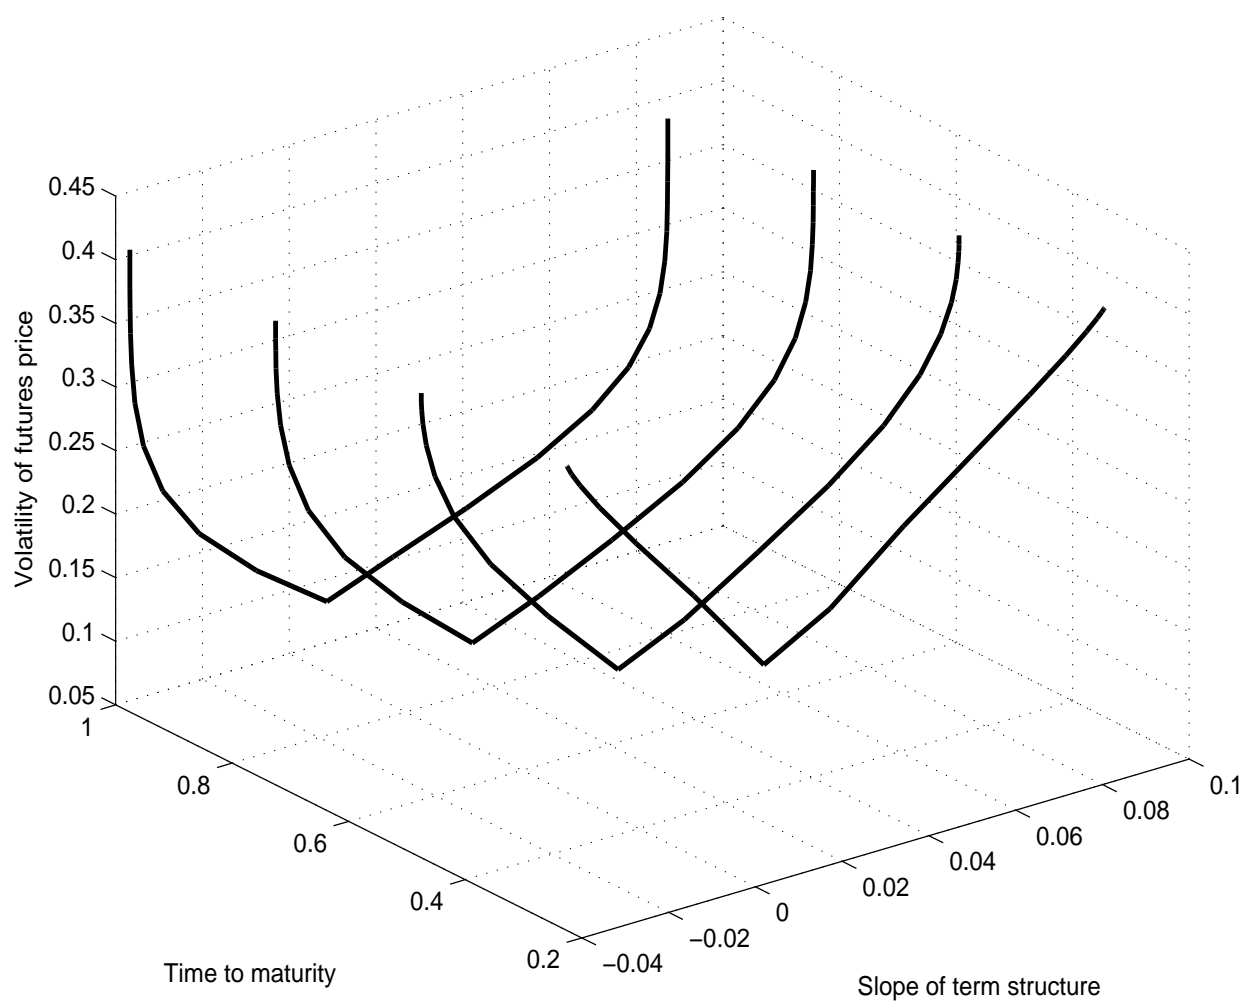


Table 1: **Conditional volatility, crude oil futures**

This Table reports results for three different regressions. The data are daily returns percentage price changes on NYMEX crude oil (CL) futures from 1983 to 2000. The specification in the both panels is the same as in Figure 4. All results are reported for times to maturity equal to one, five and ten months. All t -statistics are White-adjusted for conditional heteroscedasticity.

	1 Month	5 Months	10 Months
1. $ R(t, T) = \alpha_T + \beta_T \ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] + \varepsilon_T(t)$			
β_T (t -stat)	-0.0554 (-2.62)	-0.0455 (-2.89)	-0.0345 (-2.63)
R^2	0.0098	0.0127	0.0092
2. $ R(t, T) = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^- + \varepsilon_T(t)$			
$\beta_{1,T}$ (t -stat)	0.2517 (5.63)	0.1123 (3.02)	0.0691 (2.05)
$\beta_{2,T}$ (t -stat)	-0.1802 (-4.46)	-0.1100 (-3.87)	-0.0797 (-3.25)
R^2	0.0580	0.0381	0.0236

Table 2: **Conditional volatility, heating oil futures**

This table reports results for three different regressions. The data are daily percentage price changes on NYMEX heating oil (HO) futures from 1979 to 2000. See the caption to Table 1.

	1 Month	5 Months	10 Months
1. $ R(t, T) = \alpha_T + \beta_T \ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] + \varepsilon_T(t)$			
β_T (t -stat)	-0.0711 (-5.97)	-0.0309 (-3.51)	-0.0359 (-3.89)
R^2	0.0388	0.0140	0.0189
2. $ R(t, T) = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^- + \varepsilon_T(t)$			
$\beta_{1,T}$ (t -stat)	0.0523 (4.54)	0.0300 (3.21)	0.0192 (2.14)
$\beta_{2,T}$ (t -stat)	-0.1450 (-10.8)	-0.0646 (-4.65)	-0.0632 (-4.27)
R^2	0.0722	0.0473	0.0286

Table 3: **Conditional volatility, unleaded gasoline futures**

This Table reports results for three different regressions. The data are daily returns percentage price changes on NYMEX unleaded gasoline (HU) futures from 1985 to 2000. See the caption to Table 1.

	1 Month	5 Months	10 Months
1. $ R(t, T) = \alpha_T + \beta_T \ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] + \varepsilon_T(t)$			
β_T (t -stat)	-0.0337 (-3.89)	-0.0220 (-3.43)	-0.0222 (-2.67)
R^2	0.0107	0.0084	0.0094
2. $ R(t, T) = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^- + \varepsilon_T(t)$			
$\beta_{1,T}$ (t -stat)	0.0442 (3.45)	0.0036 (1.38)	0.0009 (1.09)
$\beta_{2,T}$ (t -stat)	-0.0795 (-6.01)	-0.0378 (-3.37)	-0.0476 (-2.76)
R^2	0.0244	0.0114	0.0161

Table 4: **Parameter estimates, crude oil futures**

This table reports our parameter values. We use a two step SMM procedure to estimate vector of seven structural parameters $\theta \equiv \{\gamma, \mu_Y, \sigma_Y, \bar{i}, r_f, \delta, \lambda\}$. Since only five model parameters can be independently identified from the data (see Section 4), we fix δ and r and estimate the remaining five parameters. We match the unconditional properties of futures prices, specifically the historic daily return on fully collateralized three-month futures position, the unconditional volatility of daily percent price changes for futures of various maturities as well as the mean, volatility, and the 30-day autoregressive coefficient of the slope of the forward curve. We use crude oil futures prices as a benchmark. In addition, we match the expected growth rate of crude oil consumption. The standard errors are reported where applicable.

Parameter	Value	Std. Error
γ	3.9221	0.0587
\bar{i}	0.1904	0.0073
r_f	0.0200	NA
μ_Y	0.0124	0.0055
σ_Y	0.1036	0.0108
δ	0.1200	NA
λ	1.5×10^{-5}	0.0016

Table 5: **Conditional volatility, model output**

Conditional variance of crude oil futures prices are compared with the output of the model under two different econometric specifications. See the caption to Table 1.

	1 Month		5 Months		10 Months	
	Data	Model	Data	Model	Data	Model
1. $ R(t, T) = \alpha_T + \beta_T \ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] + \varepsilon_T(t)$						
β_T	-0.0554	-0.0091	-0.0455	-0.0386	-0.0345	-0.0539
2. $ R(t, T) = \alpha_T + \beta_{1,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^+ + \beta_{2,T} \left(\ln \left[\frac{P(3, t-1)}{P(1, t-1)} \right] \right)^- + \varepsilon_T(t)$						
$\beta_{1,T}$	0.2517	0.0824	0.1123	0.1409	0.0691	0.1260
$\beta_{2,T}$	-0.1802	-0.2014	-0.1100	-0.4158	-0.0797	-0.4321

Table 6: **Sensitivity of model moments to parameters**

This table presents elasticities of model moments with respect to the model parameters. The baseline parameters are given in Table 4. Each elasticity is calculated by simulating the model twice: once with a value of the parameter of interest one standard deviation below its baseline value, and once with a value one standard deviation above its baseline value. Then the change in the moment is calculated as the difference between the results from the two simulations. This difference is then divided by the change in the underlying structural parameter between the two simulations. The result is then multiplied by the absolute value of the ratio of the baseline structural parameter to the baseline moment.

	Baseline Moments	γ	\bar{i}	μ_Y	σ_Y	$\lambda(\times 10^{-3})$
Average Slope of the Forward Curve	0.0101	2.2124	-0.6924	0.0738	2.2850	0.6456
Volatility of the Slope of the Forward Curve	0.0287	1.1276	1.1893	0.0537	0.1225	0.0382
30-days Autoregressive Coefficient of the Slope of the Forward Curve	0.8026	0.0616	0.0013	0.0115	0.0537	0.0029
Long Term Crude Oil Consumption Growth	-0.007	0.0021	-0.0004	-1.8046	1.5635	-1.6785
Average Return 3-month CL Futures	1.0247×10^{-4}	1.9920	-1.0826	-0.0738	2.3424	-0.1588
Unconditional Variance 1-month CL Futures	0.0230	0.9930	-0.2396	-0.0044	1.1932	-0.0032
Unconditional Variance 5-month CL Futures	0.0166	1.0025	-0.9266	-0.0274	1.7075	-0.0249
Unconditional Variance 10-month CL Futures	0.0121	1.0447	-1.6739	-0.0613	2.2210	-0.0467